Draft version December 20, 2016
Preprint typeset using LaTeX style emulateapj v. 5/2/11

arXiv:1612.05244v2 [astro-ph.IM] 19 Dec 2016

PAN-STARRS PIXEL ANALYSIS : SOURCE DETECTION AND CHARACTERIZATION

Eugene A. Magnier,[1] W. E. Sweeney,[1] K. C. Chambers,[1] H. A. Flewelling,[1] M. E. Huber,[1] P. A. Price,[2] C. Z. Waters,[1] L. Denneau,[1] P. Draper,[3] R. Jedicke,[1] K. W. Hodapp,[1] N. Kaiser,[1] R.-P. Kudritzki,[1] N. Metcalfe,[3] C. W. Stubbs,[2] R. J. Wainscoat[1]

Draft version December 20, 2016

ABSTRACT

Over 3 billion astronomical objects have been detected in the more than 22 million orthogonal transfer CCD images obtained as part of the Pan-STARRS 1 3π survey. Over 85 billion instances of those objects have been automatically detected and characterized by the Pan-STARRS Image Processing Pipeline photometry software, `psphot`. This fast, automatic, and reliable software was developed for the Pan-STARRS project, but is easily adaptable to images from other telescopes. We describe the analysis of the astronomical objects by `psphot` in general as well as for the specific case of the 3rd processing version used for the first public release of the Pan-STARRS 3π survey data.

Subject headings: Surveys:Pan-STARRS 1

1. INTRODUCTION

This is the fourth in a series of seven papers describing the Pan-STARRS1 Surveys, the data reduction techiques and the resulting data products. This paper (Paper IV) describes the details of the source detection and photometry, including point-spread-function and extended source fitting models, and the techniques for "forced" photometry measurements.

Chambers et al. (2017, Paper I) provides an overview of the Pan-STARRS System, the design and execution of the Surveys, the resulting image and catalog data products, a discussion of the overall data quality and basic characteristics, and a brief summary of important results.

Magnier et al. (2017, Paper II) describes how the various data processing stages are organised and implemented in the Imaging Processing Pipeline (IPP), including details of the the processing database which is a critical element in the IPP infrastructure .

Waters et al. (2016, Paper III) describes the details of the pixel processing algorithms, including detrending, warping, and adding (to create stacked images) and subtracting (to create difference images) and resulting image products and their properties.

Magnier et al. (2016, Paper V) describes the final calibration process, and the resulting photometric and astrometric quality.

Flewelling et al. (2016, Paper VI) describes the details of the resulting catalog data and its organization in the Pan-STARRS database. Huber et al. (2017, Paper VII) describes the Medium Deep Survey in detail, including the unique issues and data products specific to that survey. The Medium Deep Survey is not part of Data Release 1. (DR1)

The Pan-STARRS1 filters and photometric system have already been described in detail in Tonry et al. (2012).

Note: These papers are being placed on arXiv.org to provide crucial support information at the time of the public release of Data Release 1 (DR1). We expect the arXiv versions to be updated prior to submission to the Astrophysical Journal in January 2017. Feedback and suggestions for additional information from early users of the data products are welcome during the submission and refereeing process.

2. BACKGROUND

The photometric and astrometric precision goals for the Pan-STARRS 1 surveys were quite stringent: photometric accuracy of 10 millimagnitudes, relative astrometric accuracy of 10 milliarcseconds and absolute astrometric accuracy of 100 milliarcseconds with respect to the ICRS reference stars.

An additional constraint on the Pan-STARRS analysis system comes from the high data rate. PS1 produces typically ~ 500 exposures per night, corresponding to $\sim$ 750 billion pixels of imaging data. The images range from high galactic latitudes to the Galactic bulge, so large numbers of measurable stars can be expected in much of the data. The combination of the high precision goals of the astrometric and photometric measurements and the high data rate (and a finite computing budget) mean that the process of detecting, classifying, and measuring the astronomical objects in the image data stream in a timely fashion are a significant challenge.

In order to achieve these ambitious goals, the object detection, classification, and measurement process must be both precise and efficient. Not only is it necessary to make a careful measurement of the flux of individual objects, it is also critical to characterize the image point-spread-function, and its variations across the field and from image to image. Since comparisons between images must be reliable, the measurements must be stable for both photometry and astrometry.

A variety of astronomical software packages perform the basic object detection, measurement, and classification tasks needed by the Pan-STARRS IPP. Each of these programs have their own advantages and disadvantages. Below we discuss some of the most widely used of these

[1] Institute for Astronomy, University of Hawaii, 2680 Woodlawn Drive, Honolulu HI 96822
[2] Department of Astrophysical Sciences, Princeton University, Princeton, NJ 08544, USA
[3] Department of Physics, Durham University, South Road, Durham DH1 3LE, UK
[2] Harvard-Smithsonian Center for Astrophysics, 60 Garden Street, Cambridge, MA 02138

other packages, highlighting the features of the programs which are particularly desirable, and noting aspects of the programs which are problematic for the IPP.

- DoPhot : analytical fitted model with aperture corrections. pro: well-tested, stable code. con: limited range of models, algorithm converges slowly to a PSF model, limited tests of PSF validity, inflexible code base, fortran (Schechter et al. 1993).
- DAOPhot : Pixel-map PSF model with analytical component. pro: well-tested, high-quality photometry. con: Difficult to use in an automated fashion, does it handle 2D variations well? (Stetson 1987).
- Sextractor : pure aperture measurement with rudimentary object subtraction. pro: fast, widely used, easy to automate. con: poor object separation in crowded regions, PSF-modeling was only in beta, not widely used at the time (Bertin & Arnouts 1996).
- galfit : detailed galaxy modeling. not a multi-object PSF analysis tool. con: does not provide a PSF model, not easily automated. very detailed results in very slow processing. only a galaxy analysis program (Peng et al. 2002).
- SDSS phot : con: tightly integrated into the SDSS software environment (Lupton et al. 2001).

When the IPP development was starting, the existing photometry packages either did not meet the level of accuracy required or were required too much human intervention to be considered for the needs of PS1. In the case of the SDSS Photo tool, the software was judged to be too tightly integrated to the architecture of SDSS to be easily re-integrated into the Pan-STARRS pipelin. A new photometry analysis package was developed using lessons learned from the existing photometry systems. In the process, the object analysis software was written using the data analysis C-code library written for the IPP, `psLib`. Components of the photometry code were integrated into the IPP's mid-level astronomy data analysis toolkit called `psModules`. The result software, 'PSPhot', can be used either as a stand-alone C program, or as a set of library functions which may be integrated into other programs

The main version of PSPhot is a stand-alone program which is run on a single image or a group of related images representing the data read from a camera in a single exposure. The images are expected to have already been detrended so that pixel values are linearly related to the flux. The gain may be specific by the configuration system, or a variance image may be supplied. A mask may also be supplied to mark good, bad, and suspect pixels. Several variants of psphot have also been used in the PS1 PV3 analysis.

The version called PSPhotStack accepts a set of images, each representing the same patch of sky in a different filter, nominally the full *grizy* filter set for the analysis of the PS1 PV3 stack images, though where insufficient data were available in a given filter, a subset of these filters was processed as a group. As discussed in detail below, the PSPhotStack analysis includes the capability of measuring forced PSF photometry in some filter images based on the position of sources detected in the other filters. It also include an option to convolve the set of images to a single, common PSF size across the filters for the purpose of fixed aperture photometry.

A second version of PSPhot used in the PV3 analysis is called PSPhotFullForce. In this version, a set of image all representing the same pixels are processed together, with the positions of sources to be analysed loaded from a supplied file. In this version the analysis, sources are not discovered – only the supplied sources are considered. PSF models are determined for each exposure and the forced PSF photometry is measured for all sources. A subset of sources may also be used to measure forced galaxy shape parameters. As described below, a grid of galaxy models are fitted based on the supplied guess model.

3. PSPHOT DESIGN GOALS

PSPhot has a number of important requirements that it must meet, and a number of design goals which we believe will help to make usable in a wide range of circumstances. The critical requirements of the Pan-STARRS IPP which drive the requirements for PSPhot:

- **10 millimagnitude photometric accuracy**. For PSPhot, this implies that the measured photometry of stellar objects must be substantially better than this 10 mmag since the photometry error per image is combined with an error in the flat-field calibration and an error in measuring the atmospheric effects. We have set a goal for PSPhot of 3mmag photometric consistency for bright stars between pairs of images obtained in photometric conditions at the same pointing, ie to remove sensitivity to flat-field errors. This goal splits the difference between the three main contributors and still allows some leeway. This requirement must be met for well-sampled images and images with only modest undersampling.
- **10 milliarcsecond astrometric accuracy**. Relative astrometric calibration depends on the consistency of the individual measurements. The measurements from PSPhot must be sufficiently representative of the true object position to enable astrometric calibration at the 10mas level. The error in the individual measurements will be folded together with the errors introduced by the optical system, the effects of seeing, and by the available reference catalogs. We have set a goal for PSPhot of 5mas consistency between the true source postion and the measured position given reasonable PSF variations under simulations. This level must be reached for images with 250 mas pixels, implying PSPhot must introduce measurement errors less than 1/50th of a pixel. The choice of 32 bit floating point data values for the source centroids places a numerical limit of 1e-7 on the accuracy of a pixel relative to the size of a chip (since a single data value is used for X or Y). For the 4800^2 GPC chips, this yields a limit of about 0.25 milliarcsecond.

The design goals for PSPhot are chosen to make the program flexible, general, and able to meet the unknown usages cases future projects may require:

- **Flexible PSF model** Different image sources require different ways of representing the PSF. Ideally, both analytical and pixel-based versions should be possible.
- **PSF spatial variation** Most images result in some spatial PSF variations at a certain level. The PSF representation should naturally incorporate 2-D variations.
- **Flexible non-PSF models** PSPhot must be able to represent PSF-like objects as well as non-PSF sources (e.g., galaxies). It must be easy to add new object models as interesting representations of sources are invented.
- **Clean code base** PSPhot should incorporate a high-degree of abstraction and encapsulation so that changes to the code structure can be performed without pulling the code apart and starting from scratch.
- **PSF validity tests** PSPhot should include the ability to choose different types of PSF models for diffent situations, or to provide the user with methods for assessing the different PSF models.
- **Careful systematic corrections** PSPhot must carefully measure and correct for the photometric and astrometric trends introduced by using analytical PSF models.
- **User Configurable** PSPhot should allow users to change the options easily and to allow different approaches to the analysis.

4. PSPHOT ANALYSIS PROCESS

4.1. *Overview*

The PSPhot analysis is divided into several major stages:

1. **Image preparation** Load data, characterize the image background, load or construct variance and mask images.
2. **Initial object detection** Smooth, find peaks, measure basic properties.
3. **PSF determination** Select PSF candidates, perform model fits, build PSF model from fits, select best PSF model class.
4. **Bright object analysis** Fit objects with PSFs, determine PSF validity, subtract PSF-like objects, fit non-PSF model(s), select best model class, subtract model.
5. **Low S/N sources** Detect low-level sources, measure properties (aperture or PSF)
6. **Extended Source Analysis** Detailed measurements relevant to galaxies and/or other extended (non-PSF) sources.
7. **Aperture corrections** Measure the curve-of-growth, spatial aperture variations, and background-error corrections.
8. **Output** Write out objects in selected format, write out difference image, variance image, etc, as selected.

PSPhot is highly configurable. Users may choose via the configuration system which of the above analyses are performed. This is useful for testing, but also allows for specialized use cases. For example, the PSF model may already be available from external information, in which case the PSF modeling stage can be skipped.

4.2. *Image Preparation*

The first step is to prepare the image for detection of the astronomical objects. We need three separate images: the measured flux (signal image), the corresponding variance image, and a mask defining which pixels are valid and which should be ignored. The signal and variance images are represented internally as 32-bit floating point values. The variance and mask images may either be provided by the user, or they may be automatically generated from the input image, based on configuration-defined values for the image gain, read-noise, saturation, and so forth. For the function-call form of the program, the flux image is provided in the API, and references to the mask and variance are provided in the configuration information. As in the stand-alone C-program, the variance and mask may be constructed automatically by PSPhot.

The mask is represented as 16-bit integer image in which a value of 0 represents a valid pixel. Each of the 16 bits define different reasons a pixel should be ignored. This allows us to optionally respect or ignore the mask depending on the circumstance. For example, in some cases, we ignore saturated pixels completely while in other circumstances, it may be useful to know the flux value of the saturated pixel. In addition, the mask pixels are used to define the pixels available during a model fit, and which should be ignored for that specific fit. The initial mask, if not supplied by the user, is constructed by default from the image by applying three rules: 1) Pixels which are above a specified saturation level are marked as saturated (configuration keyword: `SATURATE`). 2) Pixels which are below a user-defined value are considered unresponsive and masked as dead. 3) Pixels which lie outside of a user-defined coordinate window are considered non-data pixels (eg, overscan) and are marked as invalid. The valid window is defined by the configuration variables `XMIN`, `XMAX`, `YMIN`, `YMAX`.

PSPhot (and other IPP) functions understand two types of masked pixels: "bad" and "suspect". Bad pixels are those which should not be used in any operations, while suspect pixels are those for which the reported signal may be contaminated or biased, but may be useable in some contexts. For example, a pixel with poor charge transfer efficiency is likely to be too untrustworthy to use in any circumstance, while a pixel in which persistence ghosts have been subtracted might be useful for detection or even analysis of brighter sources. Table 1 lists the 16 bit values used for PS1 mask images, along with

their description (see Waters et al. 2016, for additional information).

The variance image, if not supplied is constructed by default from the flux image using the configuration supplied values of `GAIN` and `READ_NOISE` to calculate the appropriate Poisson statistics for each pixel. In this case, the image is assumed to represent the readout from a single detector, with well-defined gain and read noise characteristics. This assumption is not always valid. For example, if the input flux image is the result of an image stack with a variable number of input measurements per pixel (due to masking and dithering), the variance cannot be calculated from the signal image alone. It is necessary in such a case to supply a variance image which accurately represents the variance as a function of position in the image.

Some image processing steps introduce cross-correlation between pixel fluxes. An obvious case is smoothing, but geometric transformations which redistibute fractional flux between neighboring pixels also introduces cross-correlations. In the noise model, it is necessary to track the impact of the cross correlations on the per-pixel variance. In the general case, this would require a complete covariance image, consisting of the set of cross-correlated pixels for each image pixel. Since a typical smoothing or warping operation may introduce correlation between 25 - 100 neighboring pixels, the size of such a covariance image is prohibitive.

Before sources are detected in the image, a model of the background is subtracted. The image is divided into a grid of background points with a spacing of 400 pixels. Superpixels of size 800×800 pixels are used to measure the local background for each background grid point, thus over-sampling the background spatial variations by a factor of 2. In the interest of speed, 10,000 randomly selected *unmasked* pixels in these regions are sampled to determine the background. Bilinear interpolation is used to generate a full-resolution image from the grid of background points, and this image is then subtracted from the science image. The background image and the background standard deviation image are kept in memory from which the values of `SKY` and `SKY_SIGMA` are calculated for each object in the output catalog. See also the discussion in Waters et al. (2016).

4.3. *Initial Object Detection*

4.3.1. *Peak Detection*

The objects are initially detected by finding the location of local peaks in the image. The flux and variance images are smoothed with a small circularly symmetric kernel using a two-pass 1D Gaussian. The smoothed flux and variance images are combined to generate a significance image in signal-to-noise units, including correction for the covariance, if known. At this stage, the goal is only to detect the brighter sources, above a user defined S/N limit (configuration keyword: `PEAKS_NSIGMA_LIMIT`). A maximum of `PEAKS_NMAX` are found at this stage. The detection efficiency for the brighter sources is not strongly dependent on the form of this smoothing function.

The local peaks in the smoothed image are found by first detecting local peaks in each row. For each peak, the neighboring pixels are then examined and the peak is accepted or rejected depending on a set of simple rules. First, any peak which is greater than all 8 neighboring pixels is kept. Any peak which is lower than any of the 8 neighboring pixels is rejected. Any peak which has the same value as any of the other 8 pixels is kept if the pixel X and Y coordinates are greater than or equal to the other equal value pixels. This simple rule set means that a flat-topped region will maintain peaks at the maximum X and Y corners of the region.

We use the 9 pixels which include the source peak to fit for the position and position errors. We model the peak of the sources as a 2D quadratic polynomial, and use a very simple bi-quadratic fit to these pixels. We use the following function to describe the peak

$$f(x,y) = C_{00} + C_{10}x + C_{01}y + C_{11}xy + C_{20}x^2 + C_{02}y^2$$

and write the Chi-Square equation:

$$\chi^2 = \sum_{i,j} (F_{i,j} - f(x,y))^2 / \sigma^2_{i,j}$$

By approximating the error per pixel as the error on just the peak, and pulling that term out of the above equation, and recognizing that the values x,y in the 3x3 grid centered on the peak pixel have values of only 0 or 1, we can greatly simplify the chi-square equation to a square matrix equation with the following values:

$$\begin{pmatrix} 9 & 0 & 0 & 0 & 6 & 6 \\ 0 & 6 & 0 & 0 & 0 & 0 \\ 0 & 0 & 6 & 0 & 0 & 0 \\ 0 & 0 & 0 & 6 & 0 & 0 \\ 6 & 0 & 0 & 0 & 6 & 4 \\ 6 & 0 & 0 & 0 & 4 & 6 \end{pmatrix} \begin{pmatrix} C_{00} \\ C_{10} \\ C_{01} \\ C_{11} \\ C_{20} \\ C_{02} \end{pmatrix} = \begin{pmatrix} \sum F_{i,j} \\ \sum F_{i,j}x \\ \sum F_{i,j}y \\ \sum F_{i,j}xy \\ \sum F_{i,j}x^2 \\ \sum F_{i,j}y^2 \end{pmatrix}$$

Inverting the 3x3 matrix terms for C_{00}, C_{20}, and C_{02}, the location of the peak is determined from the minimum of the bi-quadratic function above, and is given by:

$$x_{min} = (C_{11}C_{01} - 2C_{02}C_{10})D^{-1} \quad (1)$$

$$y_{min} = (C_{11}C_{10} - 2C_{20}C_{01})D^{-1} \quad (2)$$

$$D = 4C_{20}C_{02} - C_{11}^2 \quad (3)$$

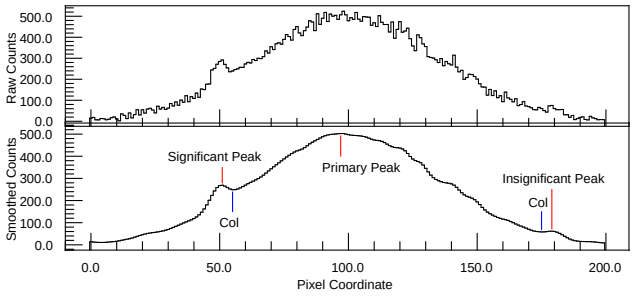

Figure 1. Illustration of peak finding and culling peaks within a footprint. Insignificant peaks within the footprint of a brighter peak are ignored in further processing.

4.3.2. *Footprints*

Table 1
PSPhot / GPC1 Mask Image Pixel Values

Mask Name	Mask Value	Description
DETECTOR	0x0001	A detector defect is present.
FLAT	0x0002	The flat field model does not calibrate the pixel reliably.
DARK	0x0004	The dark model does not calibrate the pixel reliably.
BLANK	0x0008	The pixel does not contain valid data.
CTE	0x0010	The pixel has poor charge transfer efficiency.
SAT	0x0020	The pixel is saturated.
LOW	0x0040	The pixel has a lower value than expected.
SUSPECT	0x0080	The pixel is suspected of being bad.
BURNTOOL	0x0080	The pixel contain an burntool repaired streak.
CR	0x0100	A cosmic ray is present.
SPIKE	0x0200	A diffraction spike is present.
GHOST	0x0400	An optical ghost is present.
STREAK	0x0800	A streak is present.
STARCORE	0x1000	A bright star core is present.
CONV.BAD	0x2000	The pixel is bad after convolution with a bad pixel.
CONV.POOR	0x4000	The pixel is poor after convolution with a bad pixel.
MARK	0x8000	An internal flag for temporarily marking a pixel.

The peaks detected in the image may correspond to real sources, but they may also correspond to noise fluctuations, especially in the wings of bright stars. PSPhot attempts to identify peaks which may be formally significant, but are not locally significant. It first generates a set of "footprints", contiguous collections of pixels in the smoothed significance image above the detection threshold. These regions are grown by a small amount to avoid errors on rough edges – an image of the footprints is convolved with a disk of radius 3 pixels. Peaks are assigned to the footprints in which they are contained (note by definition all peaks must be located in a footprint).

For any peak which is not the brightest peak in that footprint it is possible to reach the brightest peak by following the highest valued pixels between the two peaks. The lowest pixel along this path is the *key col* for this peak (as used in topographic descriptions of a mountain). If the key col for a given peak is less than `FOOTPRINT_CULL_NSIGMA_DELTA` (4.0) sigmas below the peak of interest, the peak is considered to be *locally insignificant* and removed from the list of possible detections (see Figure 1). In the vicinity of a saturated star, the rule is somewhat more agressive as the flat-topped or structured saturated top of a bright star may appear as multiple peaks with highly significant cols between them. However, this is an artifact of the proximity to saturation. In this regime, we require the col to also be a fixed fraction (5%) of the saturation below the peak to avoid being marked as locally insignificant.

4.3.3. *Centroid and higher-order Moments*

Once a collection of peaks has been identified, a number of basic properties of the objects related to the first and second moments are measured. Below, the second moments are used to select candidate stellar sources to be used in modeling the PSF.

In order to measure the moments, it is necessary to define an appropriate aperture in which the moments are measured. We also apply a "window function", down-weighting the pixels by a Gaussian of size σ_W which is chosen to be large compared to the PSF size, $\sigma_{\rm PSF}$. This window function reduces the noise of the measurement of the first and second moments by suppressing the noisy pixels at high radial distance as well as by reducing the contaminating effects of neighboring stars. The choice of the window function σ_W and the aperture is an iterative process: for a given value of σ_W, the PSF stars will have a measured value of $\sigma_{\rm PSF}$ which is modified by the effect of the window function. In addition, depending on the size of the window function compared to the true PSF size, the measured value of the PSF size, $\sigma_{\rm PSF}$, will be biased high or low depending on the signal-to-noise of the object.

These effects are illustrated in Figure 2 using simulated data. An image was generated with a PSF model matching the radial profile of the PS1 PSF model with a FWHM of 1.4 arcseconds. As the window function σ_W is increased, the measured FWHM for the bright simulated stars rises to meet the truth value. For small values of σ_W, fainter stars are biased to low measured values of the FWHM. For large values of σ_W, the faint stars are biased to higher values and the scatter increases.

In a real image, we do not know the true value of the PSF size. If we simply choose a very large window function and rely on bright stars, our estimate of the PSF size will be quite noisy. Compounding this problem are the two additional facts that (1) we do not know which are the real stars (as opposed to bright galaxies or possible image artifacts) and (2) the brighter stars are themselves subject to additional biases due to saturation and other non-linear effects (c.f., "the Brighter-Fatter" effect, REF). To make a robust choice for the window function $sigma_w$, we choose a value such that the measured value of $\sigma_{\rm PSF}$ is 65% of σ_w. The resulting second moment values are biased somewhat low (∼75% of the truth value for the PS1 PSF profile), but are relatively unbiased as a function of brightness.

To choose the value of σ_W, we try values of (1, 2, 3, 4.5, 6, 9, 12, 18) pixels ∼ (0.26, 0.51, 0.77, 1.15, 1.54, 2.3, 3.1, 4.6) arcseconds. For each of these values, we then select candidate PSF stars based on the distribution of the measured $\sigma_{x,x}, \sigma_{y,y}$ values. For each test value of σ_w, we determine the ratio $f = \frac{\sigma_{x,x}+\sigma_{y,y}}{2\sigma_w}$, i.e., the ratio of the window size to the observed PSF size. We interpolate to find a value of σ_W for which f is expected to be

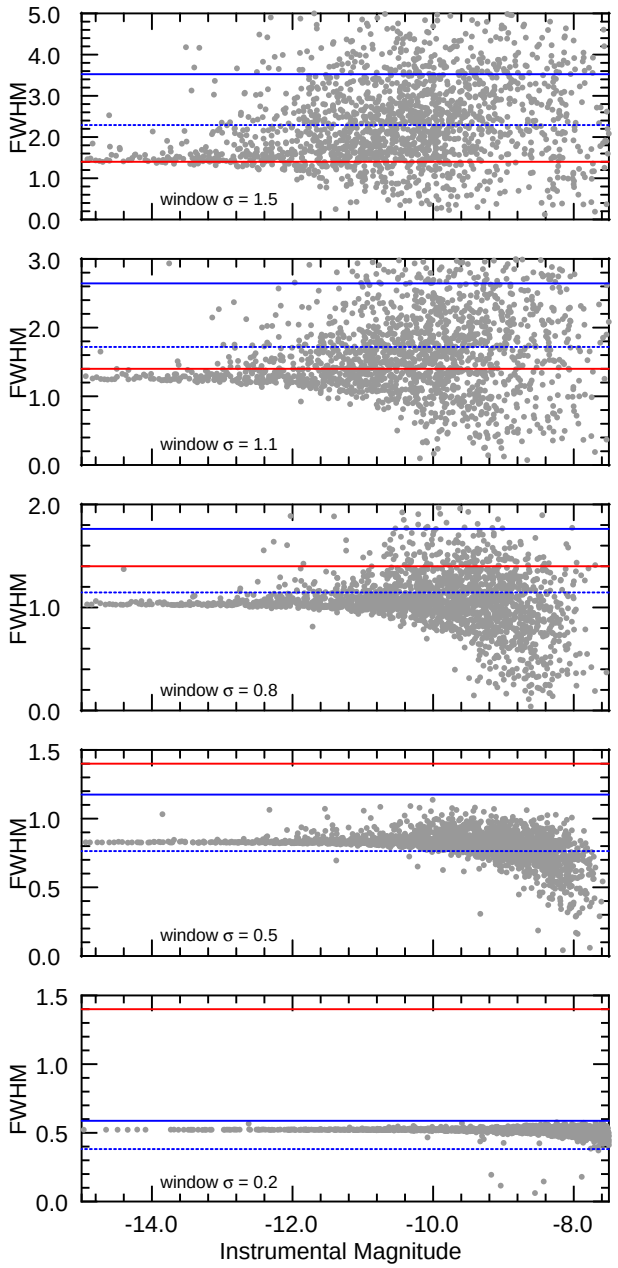

Figure 2. Example of the biases encountered when measuring the second moments. A simulated image was generated using the PS1 PSF profile. Each panel corresponds to a different value of σ_w, as marked. The solid red line is the true FWHM of the PSF used to generate the stars. The blue solid line is the FWHM of the window function ($2.35\sigma_w$). The gray dots are the FWHM derived from the measured second moments for stars in the image. The dotted blue line is the target (65% of the window function). In this example, we would choose σ_w between 0.5 and 0.8 arcseconds so the dotted blue line would match the bright end of the gray dots.

0.65. We call this value the `MOMENTS_GAUSS_SIGMA`. We use an aperture with a radius of `PSF_MOMENTS_RADIUS` = 4× `MOMENTS_GAUSS_SIGMA` to select the pixels for the measurement.

Once `PSF_MOMENTS_SIGMA` has been determined, moments are measured as defined below.

$$x_0 = \frac{1}{S}\sum_i (f_i - s_i) x_i w_i \quad (4)$$

$$y_0 = \frac{1}{S}\sum_i (f_i - s_i) y_i w_i \quad (5)$$

$$M_{xx} = \frac{1}{S}\sum_i (f_i - s_i)(x_i - x_0)^2 w_i \quad (6)$$

$$M_{xy} = \frac{1}{S}\sum_i (f_i - s_i)(x_i - x_0)(y_i - y_0) w_i \quad (7)$$

$$M_{yy} = \frac{1}{S}\sum_i (f_i - s_i)(y_i - y_0)^2 w_i \quad (8)$$

$$M_{xxx} = \frac{1}{S}\sum_i (f_i - s_i)(x_i - x_0)^3 w_i / r_i \quad (9)$$

$$M_{xxy} = \frac{1}{S}\sum_i (f_i - s_i)(x_i - x_0)^2 (y_i - y_0) w_i / r_i \quad (10)$$

$$M_{xyy} = \frac{1}{S}\sum_i (f_i - s_i)(x_i - x_0)(y_i - y_0)^2 w_i / r_i \quad (11)$$

$$M_{yyy} = \frac{1}{S}\sum_i (f_i - s_i)(y_i - y_0)^3 w_i / r_i \quad (12)$$

$$M_{xxxx} = \frac{1}{S}\sum_i (f_i - s_i)(x_i - x_0)^4 w_i / r_i^2 \quad (13)$$

$$M_{xxxy} = \frac{1}{S}\sum_i (f_i - s_i)(x_i - x_0)^3 (y_i - y_0) w_i / r_i^2 \quad (14)$$

$$M_{xxyy} = \frac{1}{S}\sum_i (f_i - s_i)(x_i - x_0)^2 (y_i - y_0)^2 w_i / r_i^2 \quad (15)$$

$$M_{xyyy} = \frac{1}{S}\sum_i (f_i - s_i)(y_i - y_0)(y_i - y_0)^3 w_i / r_i^2 \quad (16)$$

$$M_{yyyy} = \frac{1}{S}\sum_i (f_i - s_i)(y_i - y_0)^4 w_i / r_i^2 \quad (17)$$

where f_i is the flux in a pixel; s_i is the local sky value for that pixel; w_i is the value of the window function for the pixel; $S = \sum_i (f_i - s_i) w_i$ is the window-weighted sum of the source flux, used to re-normalize the moments; r_i is the radius of a pixel, $\sqrt{(x_i - x_0)^2 + (y_i - y_0)^2}$; The sum is performed over all pixels in the aperture. For the centroid calculation (x_0, y_0), the peak coordinate (see 4.3.1) is used to define the aperture and the window function; for higher order moments, the centroid is used to center the window function.

If the measured centroid coordinates (x_0, y_0) differs from the peak coordinates be a large amount (`MOMENT_RADIUS`), then the peak is identified as being of poor quality and is rejected. In both of these cases, it is likely that the 'peak' was identified in a region of flat flux distribution or many saturated or edge pixels.

In addition to the moments above, a preliminary Kron radius and flux are also calculated at this stage. In this analysis, the 1st and half-radial moments are calculated:

$$M_r = \frac{1}{S}\sum_i (f_i - s_i) r_i \quad (18)$$

$$M_h = \frac{1}{S}\sum_i (f_i - s_i)\sqrt{r_i} \quad (19)$$

Note that the window function is not applied in the calculation of these moments.

The Kron radius (Kron 1980) is defined the be 2.5× the first radial moment. The Kron flux is the sum of (sky-subtracted) pixel fluxes within the Kron radius. We also calculate the flux in two related annular apertures: the Kron inner flux is the sum of pixel values for the annulus $R_1 < r < 2.5R_1$, while the Kron outer flux is the sum of pixel values for $2.5R_1 < r < 4R_1$. The first radial moment is limited at the low and high ends by $R_{\rm min} < M_r < R_{\rm max}$ where $R_{\rm min}$ is the first radial moment of the PSF stars, or 0.75× `MOMENTS_GAUSS_SIGMA` if that cannot be determined. $R_{\rm max}$ is set to `PSF_MOMENTS_RADIUS`, the size of the moments aperture.

4.4. *PSF Determination*

4.4.1. *PSF Model vs Object Model*

PSPhot uses an analytical model to represent the shape and flux of an object. An important concept within the PSPhot code is the distinction between a model which describes an object on an image and a model with describes the point-spread-function (PSF) across an image.

Any object in an image may be represented by some analytical model, for example, a 2-D elliptical Gaussian:

$$f(x,y) = I_o e^{-z} + S \quad (20)$$

$$z = \frac{x^2}{2\sigma_x^2} + \frac{y^2}{2\sigma_y^2} + \sigma_{\rm xy} xy \quad (21)$$

$$x = x_{\rm ccd} - x_o \quad (22)$$

$$y = y_{\rm ccd} - y_o \quad (23)$$

The object model will have a variety of model parameters, in this case the centroid coordinates (x_o, y_o), the elliptical shape parameters $(\sigma_x, \sigma_y, \sigma_{\rm xy})$, the model normalization (I_o) and the local value of the background (S). A specific object will have a particular set of values for these different parameters.

The point-spread-function (PSF) of an image describes the shape of all unresolved objects in the image. In a typical image, the shape of point sources is not well described by a single function. Instead, the shape will vary as a function of position in the image. The PSF model therefore must describe the parameter variation as a function of the position of the object on the image. Note that the object model consists of a certain number of parameters which are defined by the PSF model, and another set of parameters which are independent from object to object. For the case of the elliptical Gaussian model, the PSF parameters would be the shape terms $(\sigma_x, \sigma_y, \sigma_{\rm xy})$ while the independent parameters would be the centroid, normalization and local sky values (x_o, y_o, I_o, S). Thus these parameters are each a function of the object centroid coordinates:

$$\sigma_x = f_1(x,y) \quad (24)$$

$$\sigma_y = f_2(x,y) \quad (25)$$

$$\sigma_{xy} = f_3(x,y) \quad (26)$$

$$(27)$$

PSPhot represents the variation in the PSF parameters as a function of position in the image in two possible ways, specified by the configuration. The first option is to use a 2-D polynomial which is fitted to the measured parameter values across the image. The second option is to use a grid of values which are measured for objects within a subregion of the image. In the latter case, the value at a specific coordinate in the image is determined by interpolation between the nearest grid points. The order of the polynomial or the sampling size of the grid is dynamically determined depending on the number of available of PSF stars. In the case of the PV3 analysis, the grid of values was used, with a maximum of 6×6 samples per GPC1 chip image. For the earlier PV2 analysis, the maximum grid sampling was 3×3 per GPC1 chip image. For the PV1 analysis, the polynomial representation was used, with up to 3rd order terms. The higher order representation was used for PV3 in order to follow some of the observed PSF variations in the images

PSPhot uses a single structure to represent the object model and another structure to represent the PSF model. The object model structure consists of the collection of measured object model parameters, carried as a `psLib` vector (`psVector`) along with an equal-length vector with the parameter errors. The structure also includes an integer giving the identifier of the model used in the particular case, as well as model fit statistics such as the Chi-Square of the fit and the magnitude representation of the ratio between the model flux and an aperture flux (see below for more details on this value).

The PSPhot representation of the PSF consists of an array of polynomials, each representing the variation in the object model PSF parameters (`psArray` of `psPolynomial2D`). The PSF model structure also includes the same integer used to identify which model corresponds to particular instance of the PSF. At the moment, the number of PSF parameters is a fixed number (4) fewer than the number of parameters of the corresponding object model. For example, the elliptical Gaussian model uses 7 parameters to represent the object and 3 for the PSF model.

PSPhot is written so that the object detection, measurement, and classification code does not depend on the specific form of the available object model functions. Access to the characteristics of the models is provided through a simple function abstraction method. Throughout PSPhot, there are many places where it is necessary for the code to refer to an aspect of the object or PSF model. Often, these quantities are needed deep within other parts of the code. For example, when attempting to fit the pixel flux values for an object, it is necessary to generate a guess for the model parameters. Or, in order to limit the domain of the fit, it is necessary to determine an isophotal radius for a model.

In order to avoid having the code depend on the specific form of a model, the function calls needed in these types of circumstances are abstracted, and a method is provided to return the necessary function to the higher-level software. For example, each model type has its own function to define an initial guess for the model, or a function to determine the radius for a given flux level. These are then registered as part of the model function code. Another function is then used to return the appropriate function for a specific model type. For example, the `psModelLookup_GetFunction` will return the

`psModelLookup` function for a given model type. This mechanism makes it very easy to add new model functions into the PSPhot code base. To add a new model function, the programmer simply defines a new model name (a string), the set of all necessary model lookup functions, and places the reference to the model code at the appropriate location in the psModelInit.c routine.

When a new model is provided to PSPhot, it is not necessary to specify the intended use of the object model function (ie, PSF-like object, galaxy, comet, etc). Any model can be used for the PSF model, or to describe the flux distributions of the non-PSF objects. The code currently uses a fixed translation between the object model parameters and the PSF model parameters. It also defines a specific order for the 4 independent parameters.

4.4.2. *PSF Candidate Object Selection*

The first stage of determining the PSF model for an image is to identify a collection of objects in the image which are *likely* to be PSF-like. PSPhot uses the object moments to make the initial guess at a collection of PSF-like objects. At this point, the program has measured the second order moments for all objects identified by their peaks, as well as an approximate signal-to-noise ratio. All objects with a S/N ratio greater than a user-defined parameter (`PSF_SHAPE_NSIGMA` = 20.0) are selected by PSPhot, though objects which have more than a certain number of saturated pixels are excluded at this stage. PSPhot then examines the 2-D plane of σ_x, σ_y in search of a concentrated clump of objects (see Figure 3). To do this, it constructs an artificial image with pixels representing the value of σ_x, σ_y, using a user-defined scale for the size of a pixel in this artificial image (note that the units of the σ_x, σ_y plane are the size of the second-moment in pixels in the original image). A typical value for the bin size is approximately 0.1 image pixels. The binned σ_x, σ_y plane is then examined to find a peak which has a significance greater than XXX. Unless the image is extremely sparse, such a peak will be well-defined and should represent the objects which are all very similar in shape. Other objects in the image will tend to land in very different locations, failing to produce a single peak. To avoid detecting a peak from the unresolved cosmic rays, objects which have second-moments very close to 0 are ignored. The only danger is if the PSF is very small and too many of these objects are rejected as cosmic rays.

Once a peak has been detected in this plane, the centroid and second moments of this peak are measured. All objects which land within XXX σ of this centroid are selected as likely PSF-like objects in the image.

4.4.3. *PSF Candidate Object Model Fits*

All candidate PSF objects are then fitted with the selected object model, allowing all of the parameters (PSF and independent) to vary in the fit. PSPhot uses the Levenberg-Marquardt minimization technique for the non-linear fitting. Non-linear fitting can be very computationally intensive, particularly for if the starting parameters are far from the minimization values. PSPhot uses the first and second moments to make a good guess for the centroid and shape parameters for the PSF models. Any objects which fail to converge in the fit are flagged as invalid.

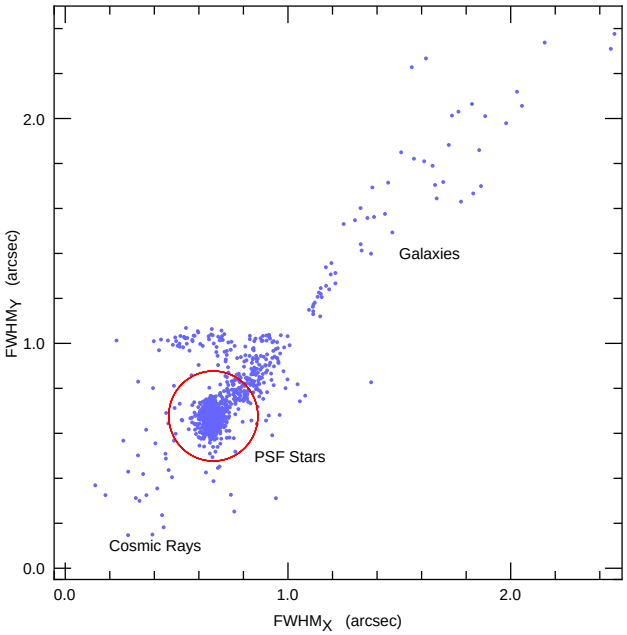

Figure 3. Illustration of PSF star selection using the FWHM derived from the second moments in $X_{\rm ccd}$ and $Y_{\rm ccd}$ directions. The dominant clump is located in this diagram. Galaxies tend to have a range of sizes and thus spread out above the stars. Cosmic rays also have a range of sizes, with one dimension smaller than the PSF. The red circle represents the PSF star candidates.

For the resulting collection of object model parameters, the PSF-dependent parameters of the models are all fitted as a function of position to a 2-D polynomial. The order of this polynomial is a user-defined parameter. The fitting process for these polynomials is iterative, and rejects the $3-\sigma$ outliers in each of three passes. This fitting technique results in a robust measurement of the variation of the PSF model parameters as a function of position without being excessively biased by individual objects which fail drastically. Objects whose model parameters are rejected by this iterative fitting technique are also marked as invalid and ignored in the later PSF model fitting stages.

All of the PSF-candidate objects are then re-fitted using the PSF model to specify the dependent model parameter values for each object. For example, in the case of the elliptical Gaussian model, the shape parameters $(\sigma_x, \sigma_y, \sigma_{xy})$ for each object are set by the coordinates of the object centroid and fixed (not allowed to vary) in the fitting procedure. The resulting fitted models are then used to determine a metric which tests the quality of the PSF model for this particular image.

The metric used by PSPhot to assess the PSF model is the scatter in the differences between the aperture and fit magnitudes for the PSF objects. The difference between the aperture and fit magnitudes (*ApResid*) is a critical parameter for any PSF modeling software which uses an analytical model to represent the flux distribution of the objects in an image. An approximate correction is measured here, with a more detailed correction measured after all object analysis is performed. The PSF model with the best consistency of the aperture correction is judged to be the best model.

4.5. *Bright Source Analysis*

4.5.1. *Fast Ensemble PSF Fitting*

Before the detailed analysis of the objects is performed, it is convenient to subtract off all of the sources, at least as well as possible at this stage. We make the assumption that all sources are PSF-like. If the centroid of the source has been determined, we use this value for its position; otherwise, we use the interpolated position of the peak. A single linear fit is used to simultaneously measure all source fluxes. Since the local sky has been subtracted, this measurement assumes the local sky is zero. We can write a single χ^2 equation for this image:

$$\chi^2 = \sum_{\text{pixels}} (F_{x,y} - \sum_{\text{sources}} A_i P[x_0, y_0])^2$$

where $F_{x,y}$ is image flux for each pixel, $P[x_0, y_0]$ is the PSF model realized at the position of source i, and A_i is the normalization for the source.

Minimizing this equation with respect to each of the A_i values results in a matrix equation:

$$M_{i,j}\bar{A}_i = \bar{F}_j$$

where $\bar{A}_i$ is the vector of A_i values, the elements of $M_{i,j}$ consist of the dot products of the unit-flux PSF for source i and source j, and $\bar{F}_j$ is the dot product of the unit-flux PSF for source j with the pixels corresponding to source j. The dot products are calculated only using pixels within the source apertures. Since most sources have no overlap with most other sources, this matrix equation results in a very sparse, mostly diagonal square matrix. The dimension is the number of sources, likely to be 1000s or 10,000s. Direct inversion of the matrix would be computationally very slow. However, an interative solution quickly yields a result with sufficient accuracy. In the iterative solution, a guess at the solution $\bar{A}$ is made assuming $M_{i,j}$ is purely diagonal; the guess is multiplied by $M_{i,j}$, and the result compared with the observed vector $\bar{F}_j$. The difference is used to modify the initial guess. This proces is repeated several times to achieve a good convergence.

Once a solution set for A_i is found, all of the objects are subtracted from the by applying these values to the unit-flux PSF.

4.5.2. *PSF Model applied to detected objects*

Once a PSF model has been selected for an image, PSPhot attempts to fit all of the detected objects, above a user-defined signal-to-noise ratio with the PSF model. For these fits, the dependent parameters are fixed by the PSF model and only the 4 independent object model parameters are allowed to vary in the fit. PSPhot again uses Levenberg-Marquardt minimization for the non-linear fitting. The objects are fitted in their S/N order, starting with the brightest and working down to the user-specified limit.

Once a solution has been achieved for an object, PSPhot attempts to judge the quality of the PSF model as a representation of the object shape. To do this, it calculates the next step of the minimization *allowing the shape parameters to vary.* This step, essentially the Gauss-Newton minimization distance from the current local minimum, should be very small if the object is well represented by the PSF, but large if the PSF is not a good representation of the object flux. The model quality is judged by the change in the two shape parameters which represent the 2D size of the object. For the case of the elliptical Gaussian, these two parameters are σ_x and σ_y. For a generic model, the shape parameters may be defined differently, but there should always be two parameters which scale the object size in two dimensions. Currently, PSPhot requires the two relevant shape parameters to be the first two dependent parameters in the list of model parameters (ie, parameters 4 & 5).

The expected distribution of the variation of the two shape parameters will be a function of the signal-to-noise of the object in question and the value of the shape parameter itself. The expected standard deviation on the shape parameter is, eg, σ_x/S/N. If the object is well-represented by the PSF, then the shape parameter values should be close to their minimization value. We can thus ask, for each object, given the measured amplitude of the Gauss-Newton step, how many standard deviations from the expected value (of 0.0) is this particular value? Objects for which the variation in the shape parameters is a large positive number of standard deviations are likely to be better represented by a larger flux distribution than the PSF (eg, a Galaxy or Comet, etc). Objects for which the variation in the shape parameters is a large negative number of standard deviations are likely to be better represented by a smaller flux distribution than the PSF (ie, a cosmic ray or other defect). A user-defined number of standard deviations is used to select these two cases, and to flag the object as a likely galaxy (really meaning 'extended') or as a likely defect.

At this stage of the analysis, PSPhot uses two additional indicators to identify good and poor PSF fits. The first of these is the signal-to-noise ratio. It is possible for the peak finding algorithm to identify peaks in locations which are not actually a normal peak. Some of these cases are in the edges of saturated, bleeding columns from bright stars, in the nearly flat halos of very bright stars, and so on. In these cases, a local peak exists, with a lower nearby sky region. However, the fitted PSF model cannot converge on the peak because it is very poorly defined (perhaps only existing in the smoothed image). The fit can either fail to converge or it can converge on a fit with very low or negative peak flux / flux normalization. PSPhot will flag any non-convergent PSF fit and any object with PSF S/N ratio lower than a user-defined cutoff. It is also useful to identify very poor fits by setting a maximum Chi-Square cutoff for objects.

As the objects are fitted to the PSF model, those which survive the exclusion stage are subtracted from the image. The subtraction process modifies the image pixels (removing the fitted flux, though not the locally fitted background) but does not modify the mask or the variance images. The signal-to-noise ratio in the image after subtraction represents the significance of the remaining flux. If the subtractions are sufficiently accurate models of the PSF flux distribution, the remaining flux should be below 1 σ significance. In practice the cores of bright stars are poorly represented and may have larger residual significance.

4.5.3. *Blended Sources*

Sources which are blended with other sources are fitted together as a set of PSFs. A single multi-object fit is performed on all blended peaks. The resulting fits are evaluated independently and any which are determined to be PSFs are subtracted from the image.

4.5.4. *Double Sources*

Sources which are judged to be non-PSF-like are confronted with two possible alternative choices. First, the object is fitted with a double-source model. In this pass, the assumption is made that there are two neighboring sources, but the peaks are blended together, or otherwise not distinguished. The initial guess for the two peaks is made by splitting the flux of the single source in half and locating the two starting peaks at +/- 2 pixels from the original peak along the direction of the semi-major axis of the sources, as measured from the second moments. In order for the two-source model to be accepted, both sources must be judged as a valid PSF source. Otherwise, the double-PSF model is rejected and the source is fitted with the available non-PSF model or models.

4.5.5. *Source Size Assessment*

After the PSF model has been fitted to all sources, and the Kron flux has been measured for all sources, PSPhot uses these two measurements, along with some additional pixel-level analysis, to determine the size class of the object. If the object is large compared to a PSF, it is considered to be *extended* and will be fitted with a galaxy model (or possibly another type of extended source model in special cases). If the object is small compared to a PSF, it is considered to be a *cosmic ray* and masked.

Extended sources are identified as those for which the Kron magnitude is significantly brighter than the PSF magnitude when compared to a PSF star. The value $dMagKP = m_{\rm Kron} - m_{\rm PSF}$, the difference between the PSF and Kron magnitudes, is calculated for each object. The median of $dMagKP$ is calculated for the PSF stars. This median is subtracted from $dMagKP$ for each star. The result is divided by the quadrature error of the PSF and Kron magnitudes and called `extNsigma`. If `extNsigma` is larger than `PSPHOT.EXT.NSIGMA.LIMIT` (3.0), the object is considered to be extended.

Cosmic Rays are identified by a combination of the Kron magnitude and the second-moment width of the object in the minor axis direction. The second-moment in the minor axis direction is calculated from M_{xx}, M_{xy}, M_{yy} as follows:

$$M_{\rm minor} = \frac{1}{2}(M_{xx} + M_{yy}) - \frac{1}{2}\sqrt{(M_{xx} - M_{yy})^2 + 4M_{xy}^2}$$

If $M_{\rm minor} < 1.2$ pixels^2 and the instrumental Kron magnitude is < -5.5, then the object is identified as a cosmic ray and the associated pixels are masked.

4.5.6. *Non-PSF Objects*

Once every object (above the S/N cutoff) has been confronted with the PSF model, the objects which are thought to be galaxies (extended) can now be fit with appropriate models for the galaxies (or other likely extended shapes). Again, the fitting stage starts with the brightest sources (as judged by the rough S/N measured from the moments aperture) and working to a user defined S/N limit.

PSPhot will use the user-selected galaxy model to attempt the galaxy model fits. In the configuration system, the keyword `GAL_MODEL` is set to the model of interest. All suspected extended objects are fitted with the model, allowing all of the parameters to float. The initial parameter guesses are critical here to achieving convergence on the model fits in a reasonable time. The moments and the pixel flux distribution are used to make the initial parameter guess. Many of the object parameters can be accurately guessed from the first and second moments. The power-law slope can be guessed by measuring the isophotal level at two elliptical radii and comparing the ratio to that expected.

For each of the galaxy models (in fact for all object models), a function is defined which examines the fit results and determines if the fit can be consider as a success or a failure. The exact criteria for this decision will depend on the details of the model, and so this level of abstraction is needed. For example, in some case, the range of valid values for each of the parameters must be considered in the fit assessment. In other cases, we may choose to use only the parameter errors and the fit Chi-Square value.

All galaxy model fits which are successful are then subtracted from the image as is done for the successful PSF model fits. Of course, the background flux is retained, with the result that only the object is subtracted from the image. Again, the variance image is (currently) not modified.

4.6. *Faint Sources*

After a first pass through the image, in which the brighter sources above a high threshold level have been detected, measured, and subtracted, PSPhot optionally begins a second pass at the image. In this stage, the new peaks are detected on the image with the bright objects subtracted. In this pass, the peak detection process uses the variance image to test the validity of the individual peaks. All peaks with a significance greater than a user-defined minimum threshold are accepted as objects of potential interest.

The objects which are measured in this faint-object stage are clearly low significance detections. The PV3 threshold for the bright object analysis is a signal-to-noise of 20. The lower limit cutoff for the faint object analysis in PV3 is a signal-to-noise of 5.0. Objects detected in the faint object stage are fitted with the PSF model using the linear, ensemble fitting process.

4.7. *Aperture Correction Measurement*

The important concept here is that an analytical model will always fail to describe the flux of the objects at some level. In the end, all astronomical photometry is in some sense a relative measurement between two images. Whether the goal is calibration of a science image taken at one location to a standard star image at another location, or the goal is simply the repetitive photometry of the same star at the same location in the image, it is always necessary to compare the photometry between two images. If this measurement is to be consistent, then the measurement must represent the flux of the stars in the

same way regardless of the conditions under which the images were taken, at least within some range of normal image conditions. So, for example, two images with different image quality, or with different tracking and focus errors, will have different PSF models. Since an analytical model will always fail to represent the flux of the star at some level, the measured flux of the same object in the two images will be different (even assuming all other atmospheric and instrumental effects have been corrected). The amplitude of the error will by determined by how inconsistently the models represent the actual object flux. For example, if the first image PSF model flux is consistently 10% too low and the second is 5% too high, then the comparison between the two images will be in error by 15%.

Aperture photometry avoids these problems, by trading for other difficulties. In aperture photometry, if a large enough aperture is chosen, the amount of flux which is lost will be a small fraction of the total object flux. Even more importantly, as the image conditions change, the amount lost will change by an even smaller fraction, at least for a large aperture. This can be seen by the fact that the dominant variations in the image quality are in the focus, tracking and seeing. All of these errors initially affect the cores of the stellar images, rather than the wide wings. The wide wings are largely dominated by scattering in the optics and scattering in the atmosphere. The amplitude and distribution of these two scattering functions do not change significantly or quickly for a single telescope and site.

The difficulty for aperture photometry is the need to make an accurate measurement of the local background for each object. As the aperture grows, errors in the measurement of the sky flux start to become dominant. If the aperture is too small, then variation in the image quality are dominant. The brighter is the object, the smaller is the error introduced by the large size of the aperture. However, the number of very bright stars is limited in any image, and of course the brighter stars are more likely to suffer from non-linearity or saturation. PSPhot measures the aperture correction (*ApResid*) for every PSF candidate object and applies this correction to the PSF model photometry.

PSPhot allows a collection of PSF model functions to be tried on all PSF candidate objects. For each model test, the above corrected ApResid scatter is measured. The PSF model function with the smallest value for the ApResid scatter is then used by PSPhot as the best PSF model for this image. The number of models to be tested is specified by the configuration keyword `PSF_MODEL_N`. The configuration variables `PSF_MODEL_0`, `PSF_MODEL_1`, through `PSF_MODEL_N - 1` specify the names of the models which should be tested.

Several likely PSF model classes are available within `psphot`:

- Gaussian : $f = I_0 e^{-z}$
- Pseudo-Gaussian : $f = I_0(1 + z + \frac{1}{2}z^2 + \frac{1}{6}z^3)^{-1}$ `[PGAUSS]`
- Variable Power-Law : $f = I_0(1 + z + z^\alpha)^{-1}$ `[RGAUSS]`
- Steep Power-Law : $f = I_0(1 + \kappa z + z^{2.25})^{-1}$ `[QGAUSS]`
- PS1 Power-Law : $f = I_0(1+\kappa z+z^{1.67})^{-1}$ `[PS1_V1]`

where $z \propto r^2$ ($z = \frac{x^2}{2\sigma_x^2} + \frac{y^2}{2\sigma_y^2} + \sigma_{xy}xy$). The Pseudo-Gaussian is a Taylor expansion of the Gaussian and is used by Dophot (Schechter et al. 1993). The latter profiles are similar to the Moffat profile form (**??**), with small differences. For the PS1 GPC1 analysis, we used the `PS1_V1` model, which we found by experimentation to match well to the observed profiles generated by PS1. Figure 4 shows example radial profiles for moderately bright stars in fairly good (0.9 arcsec) and poor (2.2 arcsec) seeing. Using a fixed power-law exponent results in somewhat faster profile fitting compared to the variable power-law exponent model.

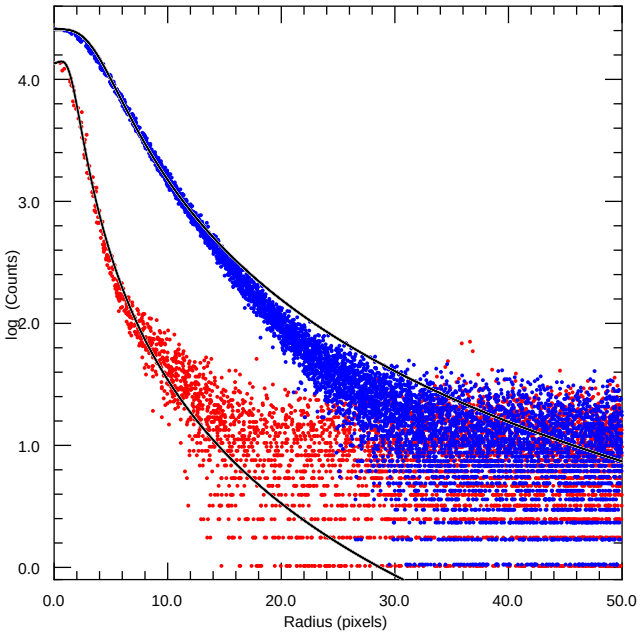

Figure 4. Radial profiles of stellar images from PS1. These two profiles illustrate the radial trend of the PS1 PSFs for a star with FWHM 0.9 arcsec (red) and 2.2 arcsec (blue). The black line shows the PSF model with radial trend of the form $(1 + \kappa r^2 + r^{3.33})^{-1}$.

4.8. *Radial Profiles*

Galaxies with regular profiles, such as elliptical galaxies and regular spiral galaxies, may be described as primarily a radial surface brightness profile, with additional structure acting as a perturbation on that profile. For many galaxies, the azimuthal shape at a given isophotal level may be described as an elliptical contour. To first order, a galaxy may be well decribed with a single elliptical contour and radial profile.

In order to facilitate the Petrosian photometry analysis below, PSPhot generates a radial profile for each suspected galaxy. This analysis starts by generating a radial profile in 24 azimuthal segments. Near the center of the galaxy, the profile is defined for radial coordinates in steps of 1 pixel, with the closest pixel values interpolated to that radial position. Further from the center,

profile is defined using the median of the pixels landing in an annular segment of size $\delta R = r \sin\theta$, rounded up to the nearest integer pixel value. The median of all pixels within a rectangular approximation to the radial wedge is used.

The resulting 24 radial profiles are subject to contamination from neighboring sources or to NAN values from masked pixels. To clean the profiles, pairs of radial profiles from opposite sides of the source are compared. Any masked values are replaced by the corresponding value in the other profile. The minimum of both profiles is the kept for both profiles. The result of this analysis is a set of profiles of the form $f_i(r_i)$. In this case, f_i is effectively the surface brightness for each radius in instrumental counts per pixel.

The surface brightness profiles are then used to define the radial contour at a specific isophotal level. This contour will be used to rescale the radial profiles into a single set of profiles normalized by the elliptical contour. This contour is defined by determining the median radius for profile bins with surface brightness in the range $F_{\rm min} + 0.1F_{\rm range}$ to $F_{\rm min} + 0.5F_{\rm range}$. The result of this analysis is a value for the radius as a function of the angle for a well-defined surface brightness regime. We then determine the elliptical shape parameters for this elliptical contour: $R_{\rm major}, R_{\rm minor}, \theta$. This ellipse is then used to redefine a single radial profile normalized by the elliptical contour:

$$\rho = \sqrt{\frac{x^2}{S_{xx}^2} + \frac{y^2}{S_{yy}^2} + xyS_{xy}}$$

The surface brightness values are sampled at a number of radial annuli, with the radii defined in the configuration (`RADIAL.ANNULAR.BINS.LOWER` & `RADIAL.ANNULAR.BINS.UPPER`). For each source, the resulting surface brightness profile is saved in the output cmf-file as an N-element value in the FITS table (`PROF_SB`). The flux at each radial position and the fill-factor (fraction of pixels used to the total possible) as also saved as equal-length vectors in the FITS table (`PROF_FLUX` and `PROF_FILL`). The values of the radial bins are saved in the cmf header (`RMIN_NN`, `RMAX_NN`).

4.9. *Petrosian Radii and Magnitudes*

Petrosian (1976) defined an adaptive aperture based on a ratio of surface brightnesses. The motivation is to define an aperture which can be determined for galaxies without significant biases as a function of distance. Since surface brightness in a resolved object is conserved, using a ratio of surface brightness to define a spatial scale results in a spatial scale which is constant regardless of galaxy distance.

To measure the Petrosian radius and flux, we start by defining a series of radial apertures with power-law spacing: $r_{i+1} = 1.25 r_i$. We calculate the surface brightness for the annulus from $r_i - r_{i+1}$ by calculating the median of the values in the range $r_i/\sqrt{1.25}$ to $r_{i+1}\sqrt{1.25}$ and dividing the the effective area of the annulus corresponding to $r_i - r_{i+1}$.

For any annulus i spanning the radii $r_{\rm min}$ to $r_{\rm max} = \beta r_{\rm min}$, the Petrosian Ratio for that annulus is defined as the ratio of the surface brightness in the annulus to the average surface brigthness within $r_{\rm max}$. The Petrosian Radius is defined to be $r_{\rm max}$ for the annulus for which the Petrosian Ratio = 0.2, i.e., the point on the galaxy radial profile at which the surface brightness is 20% of the average surface brightness at that point.

We determine the Petrosian Radius for the galaxy by quadratic interpolation between the last two of the fixed annuli with Petrosian Ratio > 0.2 and the first annulus with Petrosian Ratio < 0.2. In general, the Petrosian Ratio for a galaxy with a regular morphology (spiral or elliptical) is falling monotonically, so this interpolation is unambiguous. However, irregular galaxy morphologies, noise, and/or significant masking can cause the Petrosian Ratio to have rises as well as drops. We track the Petrosian Ratio until the value is no longer significant ($\sigma_{\rm Ratio} < 2{\rm Ratio}$). If the Petrosian Ratio drops below 0.2 for more than one radius, we choose the largest such radius.

Once the Petrosian Radius has been determined, we can now measure the Petrosian Flux : this is defined to be the total flux within an aperture corresponding to 2 $\times$ the Petrosian Radius. Using the Petrosian Flux, we can calculate two other interesting radii: R_{50} and R_{90}, the radii inside which 50% and 90% of the total Petrosian flux is contained.

4.10. *Radial Profile Wings*

We attempt to measure the radial profile of sources in order to find the radius at which the flux of the object is matches the sky. In this analysis, a series of up to 25 radial bins with power-law spacing are defined and the flux of the object in each annulus is measured. The "sky radius" is defined to be the radius at which the (robust median) flux in the annulus is within 1 σ of the local sky level. If this limit is not reached before the slope of the flux from one annulus to the next is less than a user-defined limit, then the annulus at which the slope reaches this limit is used to define the sky radius. These values are saved in the output smf / cmf files, but not sent to the PSPS. The sky radius value is used below in the calculation of the kron magnitude.

4.11. *Kron Magnitudes*

Preliminary Kron radius and flux values (Kron 1980) are calculated soon after sources are detected (Section 4.3.3). However, these preliminary values are not accurate due to the window-functions applied. After sources have been characterized and the PSF model is well-determined, the Kron parameters are re-calculated more carefully. In this version of the calculation, the image is first smoothed by Gaussian kernel with $\sigma = 1.7$ pixels, corresponding to a FWHM of 1.0″ in the PS1 stack images. Next, the Kron radius is determined in an iterative process: the first radial moment is measured using the pixels in an aperture 6$\times$ the first radial moment from the previous iteration. On the first iteration, the sky radius is used in place of the first radial moment. By default, 2 iterations are performed. The Kron radius is defined the be 2.5$\times$ the first radial moment. The Kron flux is the sum of pixel fluxes within the Kron radius. We also calculate the flux in two related annular apertures: the Kron inner flux is the sum of pixel values for the annulus $R_1 < r < 2.5R_1$, while the Kron outer flux is the sum of pixel values for $2.5R_1 < r < 4R_1$.

Two details in the calculation above should be noted. First, for faint sources, noise in the measurement of the 1st radial moment may result in an excessively small aperture for the successive calculations. The window used for the calculations is constrained to be at least the size of the aperture based on the PSF stars (Section 4.3.3). At the other extreme, noise may make the radius grow excessively, resulting in an unrealistically low effective surface brightness. The aperture is constrained to be less than a maximum value defined such that the minimum surface brightness is $1/2 times$ the effective surface brightness of a source detected at the 5σ limit.

Second, the measurement of the 1st radial moment includes a filter to reduce contamination from outlier pixels. Pairs of pixels on opposites sides of the central pixel are considered together. The geometric mean of the two fluxes is used to replace the flux values. If the object has 180° symmetry, this operation has no impact. However, if one of the two pixels is unusually high, the value will be surpressed by the matched pixel on the other side. This trick has the effect of reducing the impact of pixels which include flux from near neighbors.

4.12. *Convolved Galaxy Model Fits*

In the galaxy model fittting stage, sources which meet certain criteria are fitted with analytical models for galaxies. Three traditional analytical galaxy models are implemented in `psphot` and used in the PV3 analysis:

- Exponential profile : $f = I_0 e^{-\rho}$
- DeVaucouleur profile (de Vaucouleurs 1948): $f = I_0 e^{-\rho^{1/4}}$
- Sersic (Sérsic 1963) : $f = I_0 e^{-\rho^{1/n}}$

where ρ is a normalized radial term: $\rho = \sqrt{\frac{x^2}{R_{xx}^2} + \frac{y^2}{R_{yy}^2} + xyR_{xy}}$. The terms ($R_{xx}$, R_{yy} , R_{xy}) describe the elliptical contour and the profile scale in all three models and the coordinates x & y are determined relative to the centroids ($x, y = X_{\rm chip} - x_0, Y_{\rm chip} - y_0$). Including the normalization (I_0) and a local sky value, the Exponential and DeVaucouleur profiles have 7 free parameters and the Sersic profile has the additional free parameter of the Sersic index n.

In this stage, the galaxy model is convolved with an approximation to our best guess for the PSF model at the location of the galaxy. For the PV3 analysis, all sources detected in the 'bright source' analysis step ($S/N > 20$) were fitted with all three galaxy models, unless (a) the morphological test identified the source as a likely cosmic ray (Section 4.5.5) or (b) the peak of the PSF profile was above the saturation limit for the chip (see the discussion in Waters et al. 2016, regarding the masking of saturated pixels). Sources in the denser portions of the Galactic plane and bulge were not included in the analysis. This restriction limited the total time spent on the galaxy modeling analysis at the expense of galaxy photometry in the plane (though Kron photometry is available for those objects). The Galactic Plane region was defined by $|b| > b_{\rm min}$ where $b_{\rm min} = b_0 + r_b e^{\frac{-l^2}{2\sigma_b^2}}$. For the PV3 analysis, b_0 =20°, r_b =15°, σ_b =50°.

Before the non-linear fitting may be performed, it is necessary to determine initial values for the parameters to be fitted. For each of the three model types, the position determined from the PSF fitting analysis is used as the initial centroid x_0, y_0. A guess for the terms (R_{xx}, R_{yy} , R_{xy}) is generated based on the second moments. The guess does not attempt to use the PSF model to adjust the (R_{xx}, R_{yy} , R_{xy}) values; it was found that such a guess tended to be too small and resulted in more iterations rather than fewer. The 1st radial moment (see 4.3.3) is used to estimate the effective radius of the model based on the results of Graham & Driver (2005, Table 1). They quantive the relationships between the first radial moment used to calculated a Kron Magnitude and the effective radius for different Sersic index values, n. Since the Exponential and DeVaucouleur models are equivalent to Sersic models with $n = 1$ and 4, respectively, this work can be used to generate the initial effective radius values for all 3 model types. Once the effective radius is chosen, the second moments are used to define the aspect ratio and position angle of the elliptical contour. The Kron flux is used to generate a guess for the normalization, applying an appropriate scale factor based on the (R_{xx}, R_{yy} , R_{xy}) values, generated by integrating normalized Sersic models and determining the relationship between the central intensity and the integrated flux as a function of the Sersic index.

The PSF-convolved galaxy model fitting analysys uses the Levenberg-Marquardt minimization method to determine the best fit. In this process, the χ^2 value to be minimized is:

$$\chi^2(\bar{a}) = \sum_p \frac{1}{\sigma_p^2} \left[I_p - M_p(\bar{a}) \otimes {\rm PSF} \right]^2$$

where I_p represents the pixel values in the image (within some aperture) and $M_p(\bar{a})$ represents the unconvolved galaxy model, a function of a number of parameters $\bar{a}$, which is then convolved with the PSF model.

We simplify this by defining:

$$f_p(a_m) = \frac{1}{\sigma_p}(I_p - M_p \otimes {\rm PSF}) \tag{28}$$

$$\tag{29}$$

To determine the minimization, we need the gradient and laplacian of χ^2 with respect to the model parameters, a_m:

$$\chi^2(\bar{a}) = \sum_p f_p^2 \tag{30}$$

$$2\nabla\chi^2 = \sum_p f_p \frac{\partial f_p}{\partial a_m} \tag{31}$$

$$\nabla^2\chi^2 \sim H_{m,n} \tag{32}$$

$$2H_{m,n} = \sum_p \frac{\partial f_p}{\partial a_m}\frac{\partial f_p}{\partial a_n} \tag{33}$$

where we have approximated the Laplacian with the Hessian matrix, $H_{m,n}$ by dropping the second-derivatives (which are assumed to be a small perturbation). Since

$$\frac{\partial f_p}{\partial a_m} = -\frac{1}{\sigma_p}\frac{\partial M_p \otimes {\rm PSF}}{\partial a_m}$$

and since the order of the derivative and convolution may be exchanged, we can write these in terms of the convolved image of the model and the convolved images of the derivatives of the model M_p with respect to the model parameters, a_m:

$$\mathcal{M}_p = M_p \otimes \mathrm{PSF} \tag{34}$$

$$\mathcal{M}'_{p,m} = \frac{\partial M_p}{\partial a_m} \otimes \mathrm{PSF} \tag{35}$$

$$2\nabla\chi^2 = -\sum_p \frac{I_p - \mathcal{M}_p}{\sigma_p} \mathcal{M}'_{p,m} \tag{36}$$

$$2H_{m,n} = \sum_p \frac{1}{\sigma_p^2} \mathcal{M}'_{p,m} \mathcal{M}'_{p,n} \tag{37}$$

The gradient vector and Hessian matrix are used in the Levenberg-Marquardt minimization analysis using the standard techinique of determining a step from the current set of model parameters to a new set by solving the matrix equation:

$$(1 + \lambda_{m,n}) H_{m,n} = \delta\nabla\chi^2$$

where $\lambda_{m,n}$ is zero for $m \neq n$ and for $m = n$ set to be large when the last iteration produced a large change in the parameters compared to the local-linear expectation and small when the last change was small. The iteration ends when the change in the parameters is small and/or the change in the χ^2 value is small.

In the analysis, convolved galaxy fit, the galaxy model image and the model derivative images must be convolved with the PSF at each iteration step. To save computation time, this convolution is performed using a circularly symmetric approximation of the PSF model, with the PSF model scale size set to the average of the major and minor axis direction scale size of the full PSF model, with the same radial profile term as the PSF model. The convolution is performed directly using the circular symmetry to reduce the number of multiplications performed: all points in the 2D circularly symmetric PSF model which have the same radial pixel coordinate can be evaluated in the convolution by summing up the corresponding pixels in the (galaxy model) image to be convolved before multiplying by the PSF model profile at that radial coordinate. This approximation reduces the number of multiplications by a factor of near 8 for larger radii. For the small size of the PSF model used to convolve the galaxy model images, it was found that this direct convolution was faster than using an FFT-based convolution.

For the Exponential and DeVaucouleur fits, all parameters are fitted in the non-linear minimization stage. For the Sersic model, we do not fit the index within the Levenberg-Marquardt analysis. Instead, we start with a coarse grid search over a range of possible index values, ($n = 0.5, 1.0, 1.5, 2.0, 3.0, 4.0, 5.0, 6.0$) and a range of possible values for $R_{\rm eff}$ based on the value of R_1, the first radial moment. For a given value of the Sersic index, the $R_{\rm eff}$ is related to the 1st radial moment by the scale factor specificy by Graham & Driver. We use the observed value of the 1st radial moment and try $R_{\rm eff}$ values of a factor of (0.8, 0.9, 1.0, 1.12, 1.25) times the value predicted by the Graham and Driver equation. For each of these steps, the aspect ratio and position angle are held constant and the normalization is determined to minimize the χ^2.

We next perform 3 Levenberg-Marquardt minimization fits allowing the shape parameters (R_{xx}, R_{yy} , R_{xy}) and the normalization to be fitted, holding the centroid (x_0, y_0), Sersic index n, and sky constant. In these fits, the index n is set to the minimum value previously calculated as well as values halfway to the next, and previous, values in the grid above. E.g., if the minimum fitted index value is 3.0, then the LMM fits are performed using n = 2.5, 3.0, 3.5. The resulting χ^2 values are then used to perform quadratid interpolation to find the index n which produces the locally minium χ^2 value. Finally, this best-fit index value is held constant while Levenberg-Marquardt minimization is used to find the best fit values of all other parameters.

4.13. *Convolved Radial Aperture Photometry*

For some science goals, a well-measured color of a galaxy is more important than an accurate total magnitude. In the case of PS1, the image quality variations for stacks of different filters presents a serious challenge for the determination of precise colors. PSPhot determines a set of PSF-matched radial aperture flux measurements in order to minimize the impact of the stack image quality variations.

In PSPhotStack, the stack analysis version of PSPhot, the 5 filter images are processed together. After the PSF models have been fitted and a best set of galaxy models have been determined, three sets of radial apertures are measured. In the first set, the fluxes in the radial apertures are measured using the raw stack images. The centers of the apertures for each object across the 5 filters are fixed so that the pixels represent the equivalent portions of the same galaxy for all 5 filters. In this analysis, the best model for each object is subtracted from the image pixels for all objects excluding the object in consideration. The 'best model' is determined based on the minimum χ^2 value for the model fits.

In addition to the raw radial apertures, the stack images are each convolved with a circular Gaussian with σ chosen to yield an image with a typical FWHM of $6''$. The full set of radial apertures are again measured on these convolved images. Again, the best object models are subtracted from the image for objects not being measured. This subtraction includes the convolution to smooth the model to the effective FWHM of the convolved image. The entire procedure is then repeated with a target FWHM of $8''$.

The Pan-STARRS1 Surveys (PS1) have been made possible through contributions of the Institute for Astronomy, the University of Hawaii, the Pan-STARRS Project Office, the Max-Planck Society and its participating institutes, the Max Planck Institute for Astronomy, Heidelberg and the Max Planck Institute for Extraterrestrial Physics, Garching, The Johns Hopkins University, Durham University, the University of Edinburgh, Queen's University Belfast, the Harvard-Smithsonian Center for Astrophysics, the Las Cumbres Observatory Global Telescope Network Incorporated, the National Central University of Taiwan, the Space Telescope Science Institute, the National Aeronautics and Space Ad-

ministration under Grant No. NNX08AR22G issued through the Planetary Science Division of the NASA Science Mission Directorate, the National Science Foundation under Grant No. AST-1238877, the University of Maryland, and Eotvos Lorand University (ELTE) and the Los Alamos National Laboratory.

REFERENCES

Bertin, E. & Arnouts, S. 1996, A&AS ,, 117, 393
Chambers, K. C., Magnier, E. A., Metcalfe, N., & et al. 2017, ArXiv e-prints
de Vaucouleurs, G. 1948, Annales d'Astrophysique, 11, 247
Flewelling, H. A., Magnier, E. A., Chambers, K. C., Heasley, J. N., Holmberg, C., Huber, M. E., Sweeney, W., Waters, C. Z., Chen, T., Farrow, D., Hasinger, G., Henderson, R., Long, K. S., Metcalfe, N., Nieto-Santisteban, M. A., Norberg, P., Saglia, R. P., Szalay, A., Rest, A., Thakar, A. R., Tonry, J. L., Valenti, J., Werner, S., White, R., Denneau, L., Draper, P. W., Jedicke, R., Kudritzki, R.-P., Price, P. A., Chastel, S., McClean, B., Postman, M., & Shiao, B. 2016, ArXiv e-prints
Huber, M., TBD, A., TBD, B., & et al. 2017, ArXiv e-prints
Kron, R. G. 1980, ApJS ,, 43, 305
Lupton, R., Gunn, J. E., Ivezić, Z., Knapp, G. R., & Kent, S. 2001, in Astronomical Society of the Pacific Conference Series, Vol. 238, Astronomical Data Analysis Software and Systems X, ed. F. R. Harnden, Jr., F. A. Primini, & H. E. Payne, 269
Magnier, E. A., Schlafly, E. F., Finkbeiner, D. P., & et al. 2017, ArXiv e-prints
Magnier, E. A., Sweeney, W. E., Chambers, K. C., Flewelling, H. A., Huber, M. E., Price, P. A., Waters, C. Z., Denneau, L., Draper, P., Jedicke, R., Hodapp, K. W., Kudritzki, R.-P., Metcalfe, N., Stubbs, C. W., & Wainscoast, R. J. 2016, ArXiv e-prints
Peng, C. Y., Ho, L. C., Impey, C. D., & Rix, H.-W. 2002, AJ ,, 124, 266
Petrosian, V. 1976, ApJ ,, L209, L1
Schechter, P. L., Mateo, M., & Saha, A. 1993, PASP ,, 105, 1342
Sérsic, J. L. 1963, Boletin de la Asociacion Argentina de Astronomia La Plata Argentina, 6, 41
Stetson, P. B. 1987, PASP ,, 99, 191
Tonry, J. L., Stubbs, C. W., Lykke, K. R., Doherty, P., Shivvers, I. S., Burgett, W. S., Chambers, K. C., Hodapp, K. W., Kaiser, N., Kudritzki, R.-P., Magnier, E. A., Morgan, J. S., Price, P. A., & Wainscoat, R. J. 2012, ApJ ,, 750, 99
Waters, C. Z., Magnier, E. A., Price, P. A., Chambers, K. C., Draper, P., Flewelling, H. A., Hodapp, K. W., Huber, M. E., Jedicke, R., Kaiser, N., Kudritzki, R.-P., Lupton, R. H., Metcalfe, N., Rest, A., Sweeney, W. E., Tonry, J. L., Wainscoat, R. J., Wood-Vasey, W. M., & Builders, P. 2016, ArXiv e-prints